Avoiding Nash Inflation

Bayesian and Robust Responses to Model Uncertainty

Robert Tetlow* and Peter von zur Muehlen
Division of Research and Statistics
Board of Governors of the Federal Reserve System
Washington, DC 20551

December 2001

In his 1999 monograph *The Conquest of American Inflation* Thomas Sargent describes how a policymaker, who applies a constant-gain algorithm in estimating the Phillips curve, can fall into the grip of an induction problem: concluding on the basis of reduced-form evidence that the trade-off between inflation and output is more favorable than it actually is. This results in oscillations between periods of disinflation and reflation. The problem arises in part because the policymaker is naive about possible misspecification, his or her role in creating that misspecification, and its role in policy design. In particular, while the use of a constant-gain algorithm admits the possibility that the model may be misspecified, the policymaker does not take this into consideration when designing policy. In this paper, we relax this assumption. We derive five policy rules which treat possible misspecification in three different ways. First, the linear-quadratic Gaussian (LQG) rule exhibits the familiar pattern of escape dynamics described by Sargent. We show a rule that takes uncertainty seriously, but in a Bayesian fashion, does no better. Finally, we consider three rules that are robust in the sense of Knight. The robust rules do a worse job than the LQG approach, and sometimes a lot worse. This is so even though the induction problem faced by the policymaker provides a *prima facie* case for being robust. We conclude that there appears to be no obvious tool that can be applied mechanically to alleviate the induction problem. A corollary of this finding is that Sargent's story for the inflation of the 1970s is robust to relaxing a key assumption in the original monograph.

Keywords: uncertainty, robust control, learning, mathematical programs. JEL codes: C6, C8.

* Corresponding author. Robert.J.Tetlow@frb.gov. snail mail to mail stop 77, telephone: (202) 452-2437, facsimile: (202) 452-4843. We thank, without implication, Tom Sargent for a useful conversation and encouragement, and participants of the Society for Computational Economics conference at Yale University, June 2001. The views expressed herein are those of the authors alone and are not necessarily shared by the Board of Governors or its staff.

1. Introduction:

There are three competing explanations for the burst of inflation experienced in the late 1960s and 1970s in the United States. The first attributes the problem to the inability of the Fed to commit to a target of low inflation owing to the well-known time consistency problem.[1] According to this theory, only a change in the incentives the Fed faces will permanently alleviate the problem. One might call this the bad governance theory of inflation. The second theory suggests the Fed was victimized by a bad sequence of shocks——namely the productivity slowdown and the oil-price shocks of 1973 and 1979—that presented the Fed with an unpleasant dilemma: either force the economy into recession or tolerate a permanently higher rate of inflation.[2] We could call this the bad luck theory of inflation. Under this theory, inflation will remain low so long as the Fed's luck lasts. The third theory, laid out in Thomas Sargent's 1999 monograph *The Conquest of American Inflation* assigns the blame for the inflation of the 1970s to bad inferences from the data. Central banks in general, and the Fed in particular, created the inflation of the 1970s because they were using the wrong model. And they had the wrong model because they incorrectly inferred structural parameters from reduced-form estimates.[3]

Sargent (1999) constructs his result in two steps. First, following procedures laid out by Tinbergen (1952) and Theil (1961), a hypothetical policymaker (or authority) estimates a model of the economy using the latest available data, updating parameter estimates quarter-by-quarter. To do this, the policymaker is assumed to use a constant-gain algorithm, described below, that allows for time variation in model coefficients. Second, the estimated parameters are taken as given and the optimal policy is designed and carried out. Repeating these steps, over and over, what occurs are episodic inflations—to what is called the Nash equilibrium—followed by "escapes" to lower inflation—the Ramsey equilibrium. This arises because at the low-inflation Ramsey equilibrium, the reduced-form estimates of the Lucas supply curve show favorable trade-offs for unemployment relative to inflation. Based on this inference, the authority engenders inflation surprises. But the surprises bring about an increase in inflation—and more generally a change

1. The classic references are Kydland and Prescott (1977) and Barro and Gordon (1983). See Ireland (1999) for a modern restatement and empirical assessment of the Barro-Gordon model.
2. The "bad luck" theory was the prevailing view in the 1970s. Former Fed Vice-Chair Alan Blinder describes a variant of this view in his 1987 book. See De Long (1998) for a discourse on this subject and related themes.
3. See also DeLong (1997), Taylor (1997, 1998), Cho, Williams and Sargent (2000), and Williams (2000) for arguments along these lines.

in the time-series pattern of inflation—so that the favorable trade-off disappears. Based on this new inference, the policymaker disinflates.

The main result—episodic Nash inflations followed by disinflations to the Ramsey solution—is a consequence of two key assumptions in Sargent's set-up: first, the assumption of the constant-gain estimation which acknowledges that from the authority's perspective, the economy is subject to drift in its structural parameters; and second, the assumption that notwithstanding this acknowledgment, the policymaker takes the estimated parameters at each date as the truth, and bases policy decisions on these values. Observing that the latter assumption is at odds with the former, in this paper we relax the second assumption that the authority takes estimated parameters as given. Instead, while we retain the use of the constant-gain algorithm—or, equivalently, discounted recursive least squares—to update parameter estimates, we assume that the policymaker takes seriously the uncertainty in the estimates of these parameters.

We consider three different methods by which our policymaker might take uncertainty seriously. The first of these is *Bayesian uncertainty* in which the model parameters are assumed to have known distributions whose means and variances vary over time, and that these variances are used to design policy. The seminal references in this literature include Brainard (1967), Chow (1970) and Craine (1977). The second is *structured Knightian uncertainty* where the uncertainty is structured in the sense that it is located in one or more specific parameters of the model, but where the true values of these parameters are known only to be bounded between minimum and maximum conceivable values. Among the expositors of this approach to model uncertainty are von zur Muehlen (1982), Giannoni (2000) and Svensson (2000).[4] The third method is *unstructured Knightian uncertainty* in which the model is assumed to be misspecified in some unstructured way leading to the formulation of a game played by the central banker against a "malevolent nature". References in this strand of the literature include Caravani (1995), Hansen and Sargent (1995, 2001), Onatski and Stock (2002), Cho *et al.* (2000), Williams (2000) and Tetlow and von zur Muehlen (2001b).[5]

In considering policy responses to unstructured Knightian uncertainty this part of the paper takes one part of Sargent's work and combines it with another. In so doing we are asking if

4. There is also another, different notion of structured model uncertainty in the sense of Knight. It differs from the method used here in that the authority is assumed to choose a policy rule that maximizes the set of models for which the economy is stable. See Onatski and Stock (2002) and Tetlow and von zur Muehlen (2001b).

Sargent's proposed solution to model misspecification alleviates Sargent's induction problem, itself a manifestation of model misspecification.[6]

We compare these potential solutions of the induction problem to the linear-quadratic Gaussian control (or certainty equivalent) solution promulgated by the Tinbergen-Theil tradition.

We find that for plausible degrees of uncertainty, the Bayesian approach to uncertainty does hardly any better in terms of economic performance than does ignoring uncertainty altogether. This result is of more than academic interest. The conduct of monetary policy under uncertainty has become a subject of active interest, both in the academic literature, and more importantly within central banks. Besides the Fed, the European Central Bank, the Bank of England, the Bank of Canada, Sveriges Riksbank (Sweden), the Reserve Bank of Australia, and the Reserve Bank of New Zealand have all released working papers assessing the Bayesian approach to model uncertainty.[7] At the risk of overgeneralizing the findings, these papers generally concur with former Fed Vice-Chairman Alan Blinder's (1998) assessment that central banks should "compute the direction and magnitude of [the] optimal policy move...then do less."[8]

While the Bayesian approach to uncertainty furnishes few benefits, we also find that the robust control approach to model uncertainty can do even worse, depending in part on the technique employed and the degree of uncertainty aversion. A policymaker applying the tools of

---

5. Each of the senses in which our policymaker takes uncertainty seriously is based on the authority's ignorance of a structurally time-invariant model. Uncertainty can also be considered from "the other side" by considering the private sector's ignorance of the policy rule. Tetlow and von zur Muehlen (2001a) look at how the private sector's need to learn Taylor-type rules might affect the choice of the rule. Closer to the spirit of this paper, Bullard and Cho (2001) use the canonical New Keynesian macromodel to show that some Taylor-type rules that would be well-behaved in a world of full information allow liquidity traps to arise when private agents' expectations are based on a misspecified model.

6. That is, we take Hansen and Sargent's (1995) tools to Sargent's (1999) problem. In point of fact, however, Sargent (1999, p. 7) ascribes the particular induction problem studied here to Edmund Phelps.

7. A sampling of central bank papers on Bayesian uncertainty, in various forms, and its implications for monetary policy include Hall *et al.* (1999), Martin (1999) at the Bank of England, Schellenkens (1999) and Smets (2000) at the ECB, Shuetrim and Thompson (2000) at the Reserve Bank of Australia, Drew and Hunt (1999) at the Reserve Bank of New Zealand, Soderstrom (2000) at Sveriges Riksbank, and Srour (1999) at the Bank of Canada. Papers out of the Federal Reserve System in the United States have been legion, including Sack (2000), Rudebusch (2001), Orphanides *et al.* (2000) and Tetlow (2002).

8. The quotation, from page 11 of Blinder (1998) is actually the former Fed Vice-Chairman's characterization of lesson of Brainard, although he writes approvingly of it. Blinder's own methodology was to "use a wide variety of models and don't ever trust any one of them too much...[and to] simulate a policy on as many models as possible, throw out the outlier(s), and average the rest" (p. 12). In endnote 11 he notes that the optimal information-weighting procedure would require the use of a variance-covariance matrix. This is an example of the Bayesian approach to model uncertainty. As discussed below, robust control methods accept the first part of Blinder quotation—the part before the ellipses—but reject the second.

unstructured Knightian uncertainty exacerbates the cycles of Nash inflation followed by escape that the certainty-equivalent policy maker naively induces. The results for structured Knightian uncertainty are broadly the same, albeit less dramatic than in the case of unstructured model uncertainty. However, the precise results will depend on the range of parameter values over which the authority wants to be robust. In the empirical case we consider, the policymaker does worse than the LQG solution but one could envision boundaries on parameters that could produce different results.

In sum, our conclusion is not a happy one. A policymaker trying to avoid Sargent's induction problem, using the gamut of modern, sophisticated tools available, has no assurance of success. Indeed, in all likelihood the authority's best efforts will be met with failure. At the end of the paper, we offer a few words on the prospects for future research in alleviating the problem. Prior to that, however, immediately following this Introduction, we introduce the very simple model used in Sargent (1999) and review the methodology used there and here, to model escape dynamics from Nash equilibria. Section 3 provides a primer on Bayesian and robust control and applies it to this very simple example and presents our results. A fourth section sums up and concludes.

2. Methodology

In his 1999 monograph *The Conquest of American Inflation,* Thomas Sargent showed how the induction problem a policymaker faces can result in recurring bouts of inflation outbreaks, followed by disinflations. The induction problem is the situation the policymaker faces when he or she (wrongly) infers structural parameters from reduced-form estimates. Sargent's application, and the subsequent work by Cho *et al.* (2000) and Williams (2000), is based on the Phillips curve trade-off and follows in the line of research that begins with Lucas (1972) and Sargent (1971), culminating in Lucas's critique (1976).

In Sargent (1999), two related errors are committed. First, while the policymaker allows for the possibility that the coefficients of the Phillips curve (or the Lucas supply curve) may evolve over time, no allowance is made for such time variation to influence the way policy is formulated. If this were the only error, however, the solution might simply be a matter of using constant-term adjustments, or time-varying coefficients models to correct the problem. The second problem—the one emphasized in Lucas (1976)—is that the policymaker does not understand his or her own role, as a part of the data generating process, in determining the evolution of the reduced-form parameters. In this paper, we relax the first of these assumptions. In this section, we lay out the simplest of the models that Sargent (1999) studies, the nature of the induction problem, and the updating procedure that the policymaker is assumed to use to gather information and make inferences.

2.1 the model:

Sargent (1999) studies several models. Here we restrict our attention to the simplest—the static version of the classical Phillips curve—since doing so allows us to generate some results analytically. It is also the model that is closest in spirit to the pioneering work in Lucas (1972). The models consists of just two equations, a Lucas supply curve, and a crude policy reaction function:

$$U_t = U^* - \theta(\pi_t - \hat{\pi}_t) + \upsilon_{1t} \qquad (1)$$

$$\pi_t = \hat{\pi}_t + \upsilon_{2t} \qquad (2)$$

where U and π are the unemployment and inflation rates, respectively; and $\hat{\pi}$ is the target rate of inflation set by the central bank. Rational expectations mean that: $E\pi = \hat{\pi}$. In this simple model, $\hat{\pi}$ is taken as a control variable and so equation (2) can be interpreted as a policy rule with a control error.

The central bank's perceived model is:

$$U_t = \gamma_{0t} + \gamma_{1t}\pi_t + \varepsilon_t + \omega_t \qquad (3)$$

where ε_t is a random error, taken to be independently and identically distributed and ω_t is a mean distortion representing possible specification errors. In designing robust policy, ω_t is taken as the instrument of a hostile opponent (nature) in a Stackelberg game, as we shall outline a bit later. In the special case where $\omega_t = 0$ and the estimated parameters are taken as if they were known, policy is certainty equivalent. To this point, the presence of ω_t is the only discrepancy from what Sargent describes. The policymaker devises a rule, F, of the perceived model's parameters to guide policy:

$$\hat{\pi}_t = F(\gamma_{0t}, \gamma_{1t})Z_t \qquad (4)$$

where, in general, $Z_t = (1, U_t)'$.

2.2 discounted recursive least squares:

As in Sargent (1999), we assume that the policymaker updates estimates of the economy on a period-by-period basis. If this were done using least squares, the gain from adding periods of observations would be $1/t$, which converges on zero as $t \to \infty$. This is sensible provided one accepts that the true economy is time invariant. Under such circumstances, the first observation in a time series is just as valuable for discerning the true parameters as the most-recent observation. If, however, the policymaker wishes to entertain the possibility that the true model parameters shift over time, he or she may wish to weight recent observations more heavily than distant ones. This can be done by utilizing discounted recursive least squares:

$$\gamma_{t+1} = \gamma_t + gP_t^{-1}X_t(U_t - \gamma_t X_t) \tag{5}$$

$$P_{t+1} = P_t + g(X_t X_t' - P_t) \tag{6}$$

where $\gamma_t = \begin{bmatrix}\gamma_{0t} & \gamma_{1t}\end{bmatrix}'$, $X_t = \begin{bmatrix}1 & \pi_t\end{bmatrix}$ and $g = 1 - \rho$ with ρ being a ‘forgetting factor’ measuring the rate at which old information is discounted. Equations (5) and (6) differ from recursive least squares only in that the gain associated with each new observation is fixed at a constant, g, instead of a variable that is strictly decreasing in time. In equation (5), the vector γ_t is the *t*-dated slice of 2-by-*t* period vector time series of estimated model parameters; P_t is the “precision matrix” as of date *t* (and a part of a 2-by-2-by-*T* vector time series); and the term in parentheses is the observation error in the regression that the policymaker conducts. So equation (5) says the change in the estimated parameters is a weighted function of the observation error. In equation (6), the precision matrix is shown to evolve according to a constant proportion, g, of the discrepancy between the variance-covariance matrix of observed regression variables, XX', and the inherited precision. The constant gain has a natural Bayesian interpretation in that g can be thought of as the arrival rate of unobservable regime shifts.

2.3 policy objectives:

Following Sargent (1999), the loss function to be minimized is assumed to be quadratic in unemployment and inflation with the parameter λ measuring the disutility of unemployment relative to inflation:[9]

$$L_t = \frac{1}{2}E\sum_{t=0}^{\infty}[\lambda U_t^{\ 2} + \pi_t^2] \tag{7}$$

This loss function should be thought of as the *ex post* loss that the authority uses to compute the performance of a given strategy. In some circumstances, the authority will choose policies *ex ante* based on criteria that include uncertainty aversion, as discussed below. After the fact, however, performance will be measured by equation (7). Note that if the policymaker knew the true parameters of the model, the certainty-equivalent (linear-quadratic Gaussian, or LQG) policy would, by definition, be the best achievable outcome. The question is whether the authority's ignorance of his or her own role in generating misspecifications is sufficient for any or all of the approaches for responding to uncertainty to dominate the LQG criteria. Sargent (1999) assesses the LQG criteria for controlling the economy, effectively ignoring the model uncertainty that using the constant-gain algorithm for updating coefficient estimates explicitly admits. We relax this restriction. The ways in which we do this, and the results we obtain, are studied in the next section.

3. Controlling an Uncertain Economy

With the true economy specified by equations (1) and (2), and the most general specification of the estimated economy being equation (3), we need to specify how estimates of equation (3) generate the policy vector, F_t, with which settings of the target rate of inflation, $\hat{\pi}_t$, are chosen. Once that is done, we set the system in motion, drawing shocks, $\upsilon = \begin{bmatrix}\upsilon_{1t} & \upsilon_{2t}\end{bmatrix}$, generating observed variables, $X = \begin{bmatrix}U_t & \pi_t\end{bmatrix}$, from which regression coefficients $\gamma = \begin{bmatrix}\gamma_{0t} & \gamma_{1t}\end{bmatrix}$ are derived. Since this model is essentially static, the only state variables are the beliefs of agents. We begin with the linear-quadratic Gaussian (LQG), or certainty equivalent case. From there, we study the Bayesian prescription for parameter uncertainty: adjusting the LQG response for uncertainty as captured by the standard error of parameter estimates. Then we study robust policy from the perspective of structured and unstructured model uncertainty. In each subsection, we characterize the authority's decision rule and compute the paths for inflation and the rule parameters.

9. We follow Sargent (1999), in allowing the authority to attempt to reduce unemployment to zero even though $U^* = 5$. An alternative would have been to allow the authority to stabilize unemployment around her point estimate of U^*. It turns out, however, that such a modification makes no substantive difference to the results. The target rate of inflation is taken to be zero without loss of generality.

3.1 the optimal certainty-equivalent rule

With LQG optimization, we can write the Lagrangian for the authority's problem as:

$$\max_{\langle\pi\rangle} \quad -\frac{1}{2}\sum_{t=0}^{\infty}[\lambda U_t^2 + \pi_t^2] - \phi[U_t - \gamma_{0t} - \gamma_{1t}\pi_t] \tag{8}$$

where we note that the policymaker has taken ω_t, in equation (3), to be zero. The problem yields the following first-order conditions:

$$\pi_t - \gamma_{1t}\phi = 0 \tag{9}$$

$$\lambda U_t + \phi = 0 \tag{10}$$

from which we can write the target rate of inflation as:

$$\hat{\pi}_t^{ce} = -\lambda\gamma_{0t}\gamma_{1t}/(1 + \lambda{\gamma_{1t}}^2) \tag{11}$$

In the true economy, agents take the authority's inflation target as given and, basing inflation expectations on that target, react accordingly. Substituting $\hat{\pi}^{ce}$ into equation (1), the implied paths of unemployment and inflation (dropping time subscripts) are

$$\pi_t = -\lambda\gamma_0\gamma_1/(1 + \lambda\gamma_1^2) + \upsilon_{2t} \tag{12}$$

$$U_t = U_t^* - \theta\upsilon_{2t} + \upsilon_{1t} \tag{13}$$

Notice that using Sargent's settings of $U^* = 5$, $\lambda = 1$ and $\theta = 1$, the target and actual inflation rates will not generally be equal to zero. With, $\sigma_{\upsilon 1} = \sigma_{\upsilon 2} = 0.3$, also taken from Sargent (1999), equations (12) and (13) become the data-generating process from which the policymaker re-estimates parameters $\gamma = \begin{bmatrix}\gamma_{0t} & \gamma_{1t}\end{bmatrix}$ through time. As Sargent (1999), Cho *et al.* (2000) and Williams (2000) have shown, when the policymaker uses a fixed-gain algorithm, like the discounted recursive least squares algorithm used here, the model generates escape dynamics wherein the economy will tend towards a high-inflation outcome, called the Nash inflation equilibrium, so long as the perceived trade-off between inflation and unemployment is regarded as favorable (i.e.; when $E\theta = \gamma_{1t}$ is large). Under such conditions, the policymaker, taking γ_{1t} as predetermined, eases policy in an attempt to reduce U, is met not only by higher inflation than originally anticipated, but also by an altered trade-off for inflation. This is because the very act of pushing on the Phillips curve produces the data that reveal a smaller γ_{1t}. Once the Nash equilibrium is obtained, a sequence of shocks eventually arises that makes it worthwhile for the authority to disinflate—to "escape" from the Nash equilibrium—bringing the economy to the low-inflation, or Ramsey, equilibrium.

Using the constant-gain algorithm for updating perceived model parameters with $\rho = 0.9725$ yields the results for the LQG-controlled economy, shown in Figure 1. The Nash equilibrium inflation rate is about 5 percent, with the Ramsey solution at about zero. The Ramsey solution arises when the policymaker's updating of γ_{it} arrives on about $\{5, -1\}$ at which point the authority believes the benefits of inflating the economy are large (owing to the large absolute value of γ_{1t}). So the policymaker generates a sequence of monetary policy surprises, and inflation slowly ratchets upward. This continues until the Nash solution approaches, where γ_{1t} converges on zero and γ_{0t} is large and positive.

In fact, the economy rarely gets close to its Nash equilibrium before a chance sequence of shocks combines with the perceived steepness of the supply curve to induce a disinflationary episode. The precise timing of escapes from Nash inflation depends on the sequence of shocks. Cho *et al.* (2000) and Williams (2000) discuss the necessary conditions in some detail. Nonetheless, the pattern described in the text is clear: recurring episodes of rising inflation, reaching near the Nash equilibrium, followed by discrete disinflations; and then the process begins anew.

Figure 1
Inflation Performance and Perceived Supply Curve Parameters
(Certainty equivalent policy)

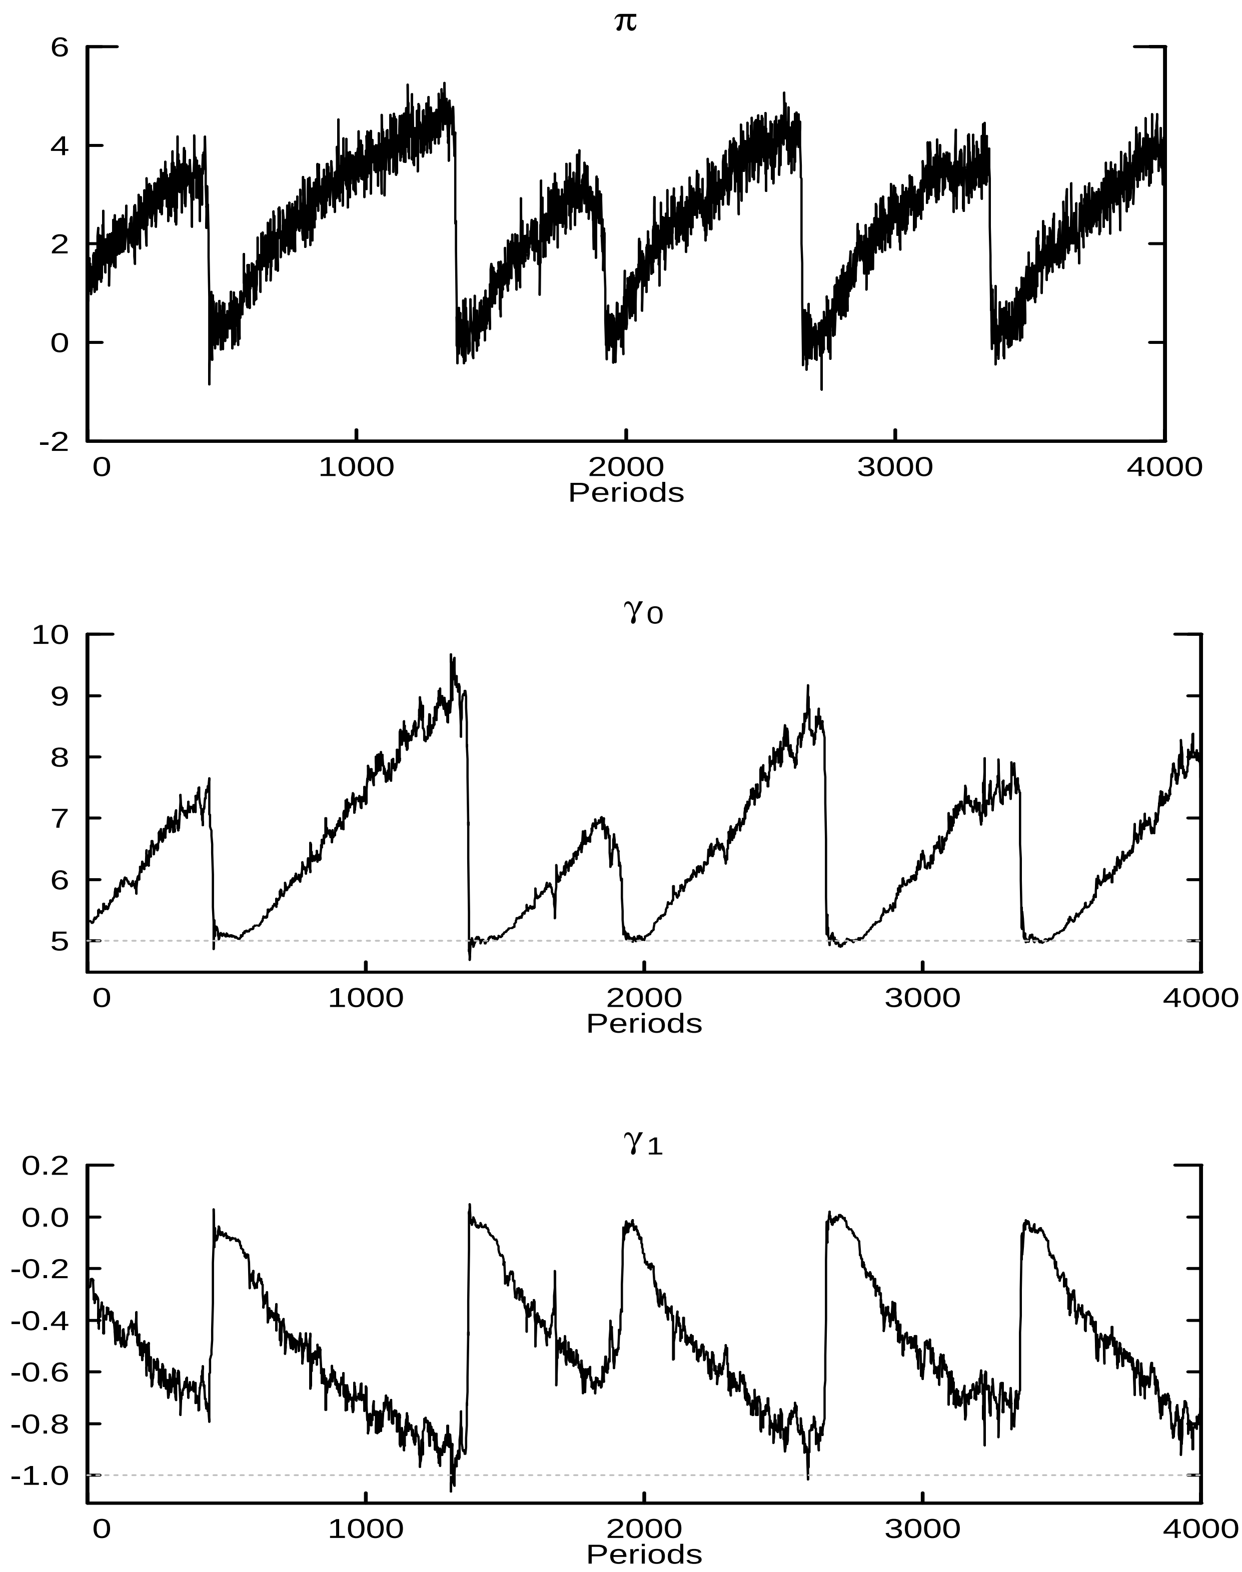

3.2 the optimal Bayesian rule

As argued above, discounted learning is based on the notion that parameters vary over time. Yet the LQG policymaker ignores this idea when setting the policy response. The Sargent characterization of policymaking seems contradictory in this sense. One possible response to the parameter variation shown in Figure 1 is to take it as a purely statistical phenomenon, much like Blinder (1998), following Brainard (1967), has advocated. In present circumstances, this amounts to respecifying the loss function as follows:

$$\max_{\langle\pi\rangle} -E\frac{1}{2}\sum_{t=0}^{\infty}[\lambda U_t^2 + \pi_t^2]$$

$$= \sum_{t=0}^{\infty}\left[E-\frac{1}{2}\pi_t^2 - \lambda\frac{1}{2}E(\gamma_{0t} + \gamma_{1t}\pi_t)^2\right] \tag{14}$$

$$= \sum_{t=0}^{\infty}\left[-\lambda\frac{1}{2}\gamma_{0t}^2 - \frac{1}{2}(1 + \lambda E\gamma_{1t}^2)\pi^2 - \lambda E\gamma_{0t}\gamma_{1t}\pi_t\right] \tag{15}$$

where we have assumed that the error terms, $\upsilon_{1,t}$ and $\upsilon_{2,t}$ are independent. This means, as before, that $E\pi_t = \hat{\pi}$. Since for this problem the model is essentially static, let us simplify the notation a bit by dropping the time subscripts where no confusion would be caused by this. The first-order condition for this problem can then be written as:

$$0 = (1 + \lambda E\gamma_1^2)\pi + \lambda E\gamma_0\gamma_1 = [1 + \lambda(\gamma_1^2 + \sigma_{\gamma 1}^2)]\pi + \lambda(\gamma_0\gamma_1 + \sigma_{\gamma 0\gamma 1}) \tag{16}$$

which implies that the policymaker’s optimal inflation target is:

$$\hat{\pi}^b = \frac{-\lambda(\gamma_0\gamma_1 + \sigma_{\gamma 0\gamma 1})}{1 + \lambda(\gamma_1^2 + \sigma_{\gamma 1}^2)} \tag{17}$$

In equation (17), the inflation target for the Bayesian policymaker differs from the LQG (certainty-equivalent) policy in equation (11) only in the presence of the terms in σ, which will vary over time. So the difference over time in the performance of the economy under the two policies (holding constant the sequence of shocks to which the economy is subjected) is in the σ and the associated differences in the vector, γ. In principle, since $\gamma^{ce} \neq \gamma^b$, and since $\sigma_{\gamma 0\gamma 1}$ could be positive or negative, the Bayesian policy could be attenuated or anti-attenuated relative to the CE policy.

Using the same parameters and the same sequence of random shocks as in Figure 1, the performance of the economy under the control of the Bayesian policymaker is shown in Figure 2 below. Even a cursory examination of this figure, relative to Figure 1, shows that there is very little to distinguish between the two. (In fact, they look identical to the unaided eye.) The facts are that despite the misspecification of the model, the time variation in the estimated parameters is insufficient to induce enough of a standard error of the estimated parameters to alter the designated policy in a significant way. In this model, the standard errors of the estimated parameters are rarely more than half the value of the estimated parameters themselves—that is, the t-statistics are at or higher than 2 nearly always, despite the misspecification. So from the perspective of the Bayesian policymaker, the (Bayesian) modeling error is not large. This is a manifestation of the self-confirming nature of the specification error the authority makes: Although the authority's model is misspecified, it is not so fundamentally misspecified as to generate a lot of variability in the reduced-form parameter estimates.

In fact, although it is not shown here, the policymaker would have to behave as if the estimated standard errors were an order of magnitude higher than what is estimated here for it to make an appreciable difference in the performance of the economy. That said, when the standard errors are as large as that, the resulting economic performance is an improvement relative to the CE policy response. That is, the Bayesian policy maker moves in the right direction, but just a trivial amount in that direction, for plausible parameterizations of the model.

There are at least some hints of this result in the literature. It is primarily in papers that work with artificial examples, such as Brainard's original 1967 paper, and more recently Söderström (1999) that uncertainty in the Bayesian sense has appeared to matter. In papers that have used real-world examples, where, for example, researchers have tried to explain the observed "sluggishness" of U.S. monetary policy, uncertainty has been insufficient as an explanation (see, e.g., Sack (2000) and Rudebusch (2001)).

Figure 2
Inflation Performance and Perceived Supply Curve Parameters
(Bayesian response to model uncertainty)

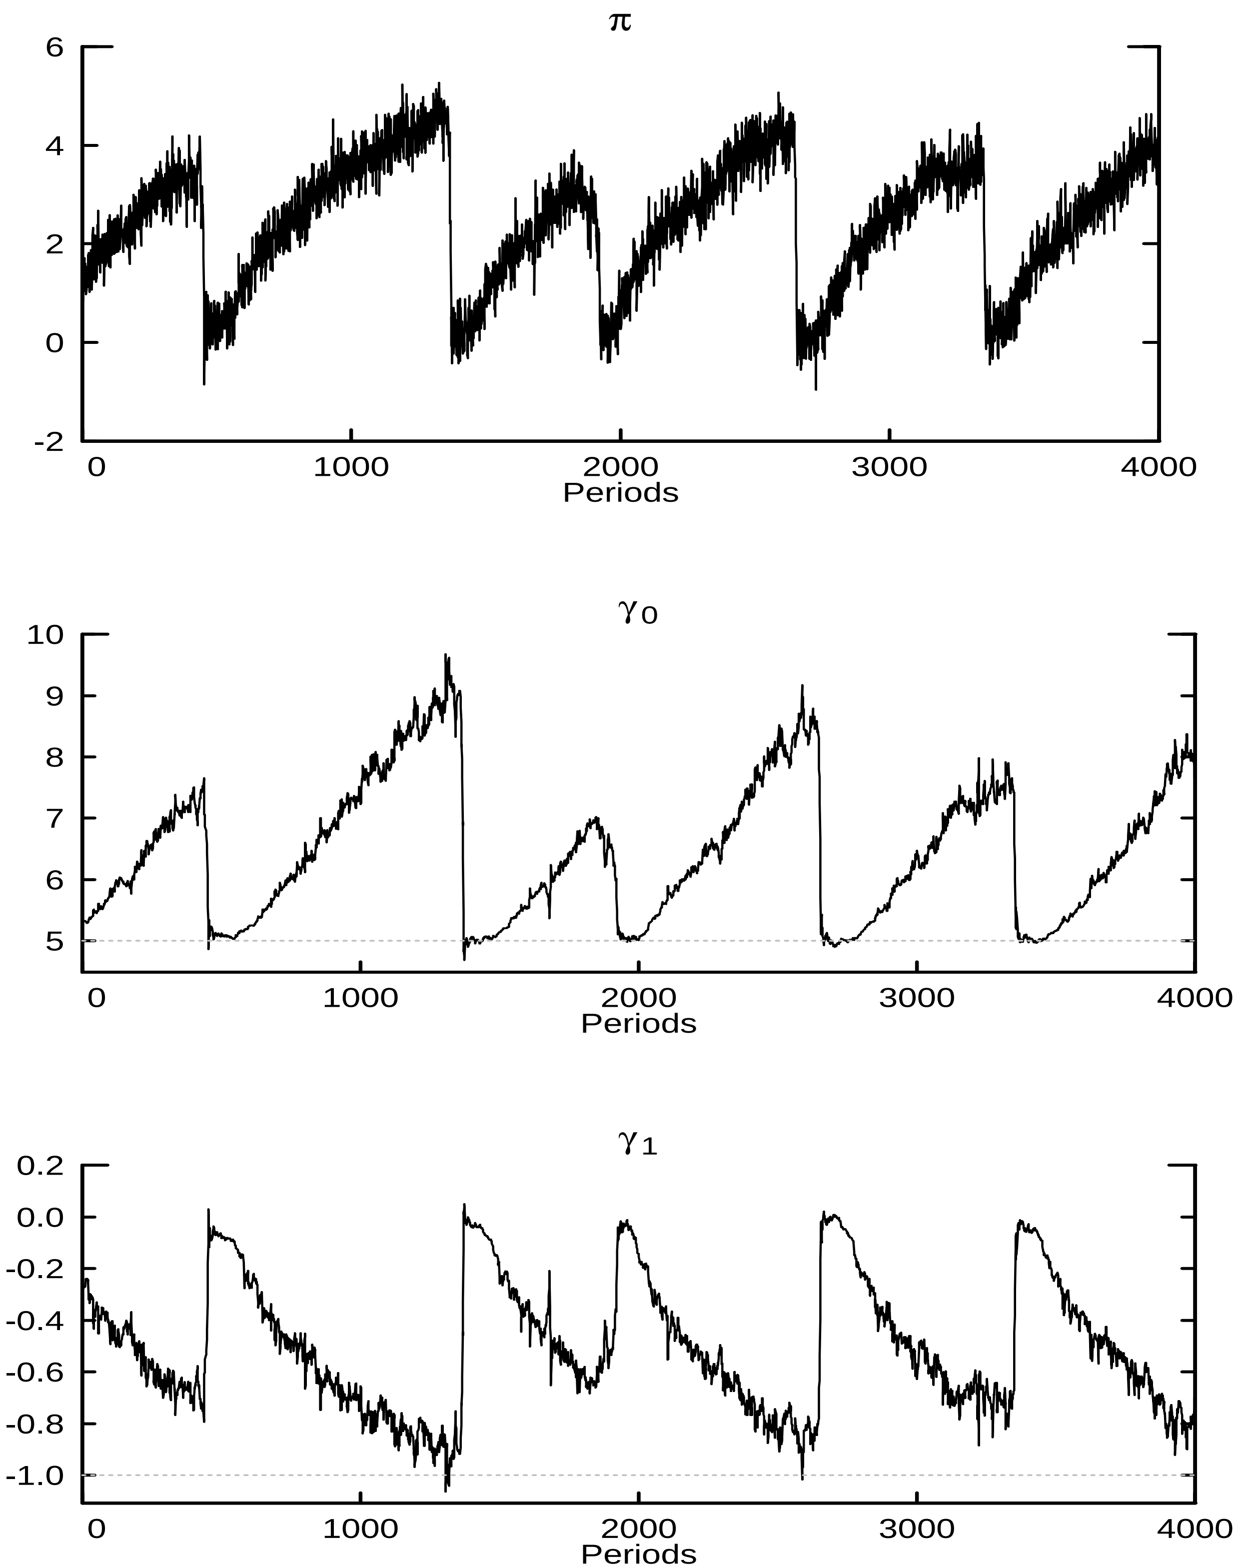

3.3 robust policy I: structured model uncertainty

In this subsection and the next one, we consider a policymaker that takes a more jaundiced view of the specification of the model. In particular, the authority is now assumed to behave as if the model were misspecified in an imprecise way. In the present subsection, we consider *structured model uncertainty*, where the authority is assumed to face Knightian uncertainty about specific parameters within a *reference model* that is otherwise taken to be approximately correct. The difference from the Bayesian concept of parameter uncertainty is two fold: First, because the authority accepts that the parameter variation stems from misspecification, as opposed to sampling error, no standard error can be assigned to the uncertainty. The best the policymaker can do is to specify a neighborhood for the uncertainty. The policymaker does this by assigning boundaries for model parameters. Second, having defined a set of models within which the true model is believed to exist, the policymaker is assumed to conduct policy to minimize the worst-case outcome within that set (see von zur Muehlen (1982), Giannoni (2000), and Svensson (2001)). This ensures that the authority has chosen a rule that will perform well for any model within the allowable set.

It is important to recognize that this set need not be large and need not include wildly improbable models. Provided that the range of models contained within the allowable set is “reasonable”, there is nothing paranoid in the authority’s response to model uncertainty.[10]

This is not the only way in which one can specify model uncertainty in the sense of Knight. In the next subsection, we shall specify another, less structured notion of Knightian uncertainty, where the assumption that the location of model misspecification is known *a priori* is dropped.

Because the authority accepts that the reference model is misspecified, there is no learning in the sense of trying to extract the true parameters from the data. Rather, the policymaker is assumed to take only the upper and lower bounds of the parameters from the lowest and highest values in historical experience:[11]

10. In fact, von zur Muehlen (1982) shows that for a linear model, the minimax solution to structured Knightian uncertainty is the same response one would obtain for a uniform distribution of parameters. Under either assumption, the only relevant cases are the boundary values for the parameters.
11. Given these bounds, we assume as before that the authority uses the same constant-gain algorithm as before for adding observations of γ_{it}.

$$\gamma_{it} \in \left[min\left\{ \gamma_{i\tau}\Big|_{\tau = 0}^{t} \right\}, max\left\{ \gamma_{i\tau}\Big|_{\tau = 0}^{t} \right\} \right] \tag{18}$$

Noting that γ_1 is generally negative, the policymaker acts to protect against worst-case scenarios of the coefficient sets: $\gamma_{0t} \in (\underline{\gamma}_0, \bar{\gamma}_0)$ and $\gamma_{1t} \in (\underline{\gamma}_1, \bar{\gamma}_1)$. The worst-case settings of these combinations are always, in a linear model, combinations of the boundaries of the two sets of coefficients.

Nature chooses $\gamma_i, i = 1, 2$ to *maximize* the loss function:

$$L_t = E\left[\frac{1}{2}\lambda(\gamma_{0t}^2 + \gamma_{1t}^2\pi_t)^2 + \pi_t^2\right] = E\frac{1}{2}[\lambda U(\gamma_0, \gamma_1)^2 + \pi_t^2] \tag{19}$$

given the monetary authority's loss minimizing choice of $\hat{\pi}$. For any choice of π, welfare is minimized if $|U(\gamma_0, \gamma_1)| = |\gamma_0 + \gamma_1\pi_t|$ is maximized. It is easily verified that for any choice of π_t, the welfare minimizing choice for nature satisfies:

$$|U(\underline{\gamma}_0, \bar{\gamma}_1)| = |U(\bar{\gamma}_0, \underline{\gamma}_1)| > |U(\underline{\gamma}_0, \underline{\gamma}_1)| = |U(\bar{\gamma}_0, \bar{\gamma}_1)| \tag{20}$$

so that the worst-case solution against which the authority must protect the economy, is either $(\underline{\gamma}_0, \bar{\gamma}_1)$ or $(\bar{\gamma}_0, \underline{\gamma}_1)$. Following von zur Muehlen (1982), the policy rule that minimizes the worst loss is:

$$\hat{\pi}_t = -\left(\frac{\underline{\gamma}_{0t} + \bar{\gamma}_{0t}}{\underline{\gamma}_{1t} + \bar{\gamma}_{1t}}\right) \tag{21}$$

that is, it is the policy that is at the midpoint of the supports of the two parameters. The policymaker's target rate of inflation will be a constant so long as the boundary values of the parameters remain constant. Just as clearly though, the precise performance of the economy under structured model uncertainty will depend critically on what determines these boundaries and the extent they vary over time if at all. Our representation of this issue is the most reasonable one we can think of, but there are others that one could devise that would have their own implications for performance.

Figure 3
Inflation Performance and Perceived Supply Curve Parameters
(Robust response to structured model uncertainty)

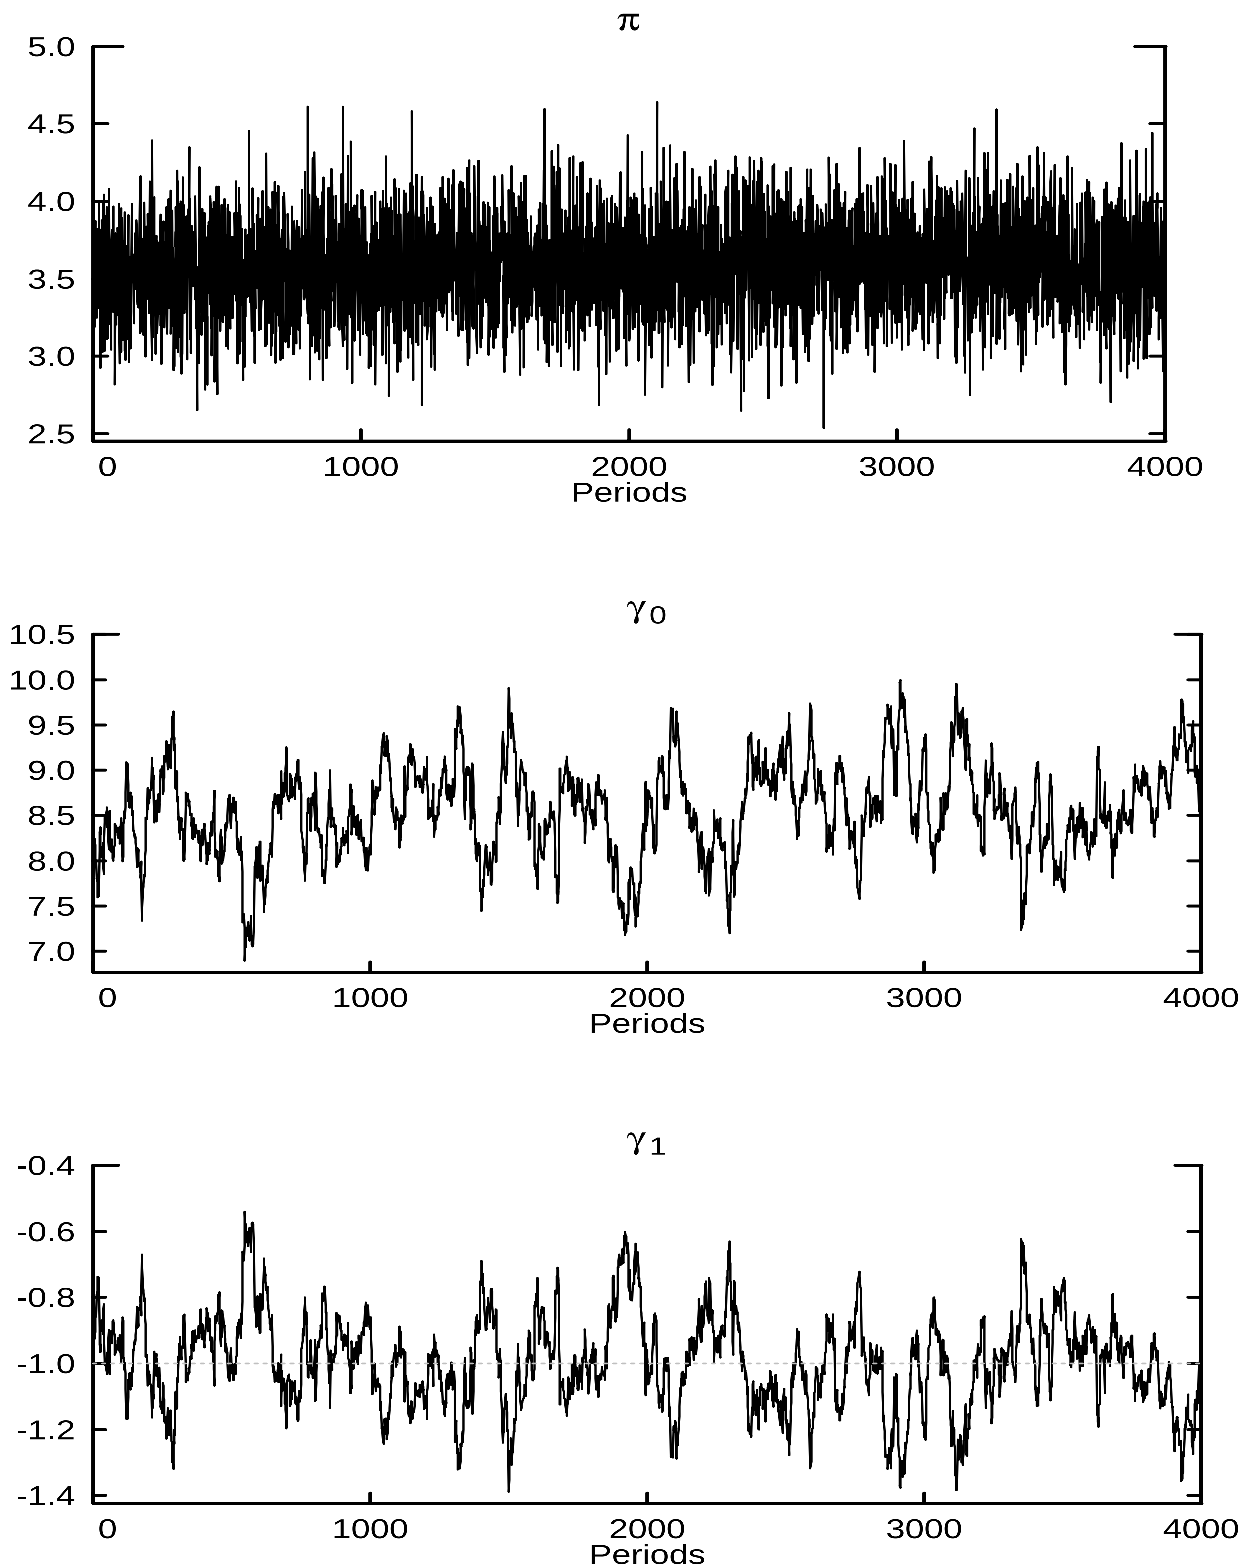

Notice that equation (21) is independent of the taste parameter λ. This was not so for the Bayesian policy maker, nor will it be so for the unstructured robust policy maker discussed below. This is because in this model worst-case parameter realizations affect welfare only through the supply curve—that is, only through their effect on inflation, whatever the choice of λ.[12]

Figure 3 shows that with this form of structured parameter uncertainty, inflation fluctuating randomly around a value of about 3-1/2 percent with the precise value being determined by the evolution of the parameter supports. We will discuss the welfare implications of this performance a bit later.

3.4 robust policy II: unstructured model uncertainty

In the previous subsection, the authority was assumed to know the location of Knightian uncertainty, but be unable to assign a probability distribution over parameters. In this subsection, we drop the assumption that the location of the Knightian uncertainty is known.That is, we assume *unstructured model uncertainty* in the sense of Knight: The uncertainty could be in the parameters, it could be in the functional form, or it could be in the perceived standard errors of the forcing shocks.

Without a location for the uncertainty, all specification errors appear to our authority as a vector of residuals. As Hansen and Sargent (1995, 2001) have argued, this results quite naturally in the policymaker acting as though he or she were the leader in a two-player Stackelberg game played against a "malevolent nature". The idea is that the misspecification is of unknown origin, but it will show up as outsized and deleterious residuals just when the policymaker attempts to exploit the model, conditional on the estimated parameters, to achieve policy objectives. The authority is best able to avoid disappointment by acting as though nature gets to choose a sequence of shocks to *maximize* the policymaker's loss, within a boundary, after the policymaker has chosen a policy rule. It follows that the policymaker will choose the rule for which the maximum loss that can be inflicted by nature is at its minimum. (See also Onatski (1999), Onatski and Stock (2002) and Tetlow and von zur Muehlen (2001b).) This is the same criterion as was used in structured model uncertainty. Getting beyond the Stackelberg game metaphor, what this set-up does is ensure that the chosen rule is optimal for the complete class of models in the allowable set, just as in the previous subsection.

12. This result will not always arise in more complicated models.

In terms of equation (3), robustness against unstructured model uncertainty is invoked by activating the expectations distortion variable—that is, the “residual,” ω_t. The assumption that the reference model is taken only to be an approximation of the truth is captured by an added constraint, equation (22) below, on the maximization of equation (8):

$$\sum_{t=0}^{\infty} \omega_t^2 \leq \eta_t^2 \qquad\qquad |\eta| < \infty \tag{22}$$

A small value for η reflects an assumption on the part of the authority that the approximation of the reference model is a close one. We use Hansen and Sargent (2001) as a guide to specify the following multiplier game:

$$\max_{\pi} \sup_{\theta > 0} \inf_{\omega} \; -\frac{1}{2}(1+\lambda\gamma_1^2)\pi^2 - \frac{1}{2}\lambda\omega^2 - \lambda\gamma_0\gamma_1\pi - \lambda\gamma_0\omega - \lambda\gamma_1\omega\pi - \frac{1}{2}\lambda\gamma_0^2 + \frac{1}{2}\mu(\omega^2 - \eta^2) \tag{23}$$

where μ is the Lagrange multiplier associated with the constraint imposed on optimization by nature’s attempt to do damage to our policymaker’s plans. In the game laid out in equation (23), the value of μ is, in some sense, a choice parameter, reflecting the extent to which the authority wishes to protect against uncertain damage.[13] Hansen and Sargent refer to μ as a preference for robustness. Nature’s influence on welfare is more limited as $\mu \to \infty$; when $\mu = \infty$ the authority chooses not to protect against model uncertainty (or equivalently there is no *ex ante* model uncertainty). As μ falls, the authority’s preference for robustness is increasing and consequently nature’s influence on policy is rising. When μ reaches λ, policy is at its most uncertainty averse; this is the H^{∞} solution; see, e.g., Whittle (1990, pp. 207-13), Caravani (1995).

Substituting for U using equation (3), the first-order conditions with respect to $\hat{\pi}$ and ω are:

$$-(1+\lambda\gamma_1^2)\pi - \lambda\gamma_1\omega - \lambda\gamma_0\gamma_1 = 0 \tag{24}$$

$$-\lambda\gamma_1\pi + (\mu - \lambda)\omega - \lambda\gamma_0 = 0 \tag{25}$$

Equation (25) clearly shows that the magnitude of μ matters in determining the outcome, and

13. We say in the text that μ is a choice parameter “in some sense” because there is the delicate, almost existential issue of whether one can choose how much to hedge against model misspecification of a given magnitude, or whether that parameter is given by tastes. Formally, the “choice” of μ is directly determined by the magnitude of uncertainty, η. This is easy to see in the present example: Solve equation (25)for μ and set $|\omega| = |\eta|$, this being nature’s best choice when $\mu > \lambda$. This gives $\mu = \lambda(1 + |\gamma_0 + \gamma_1\pi| / \eta)$. Looking at this, we can see that as $\eta \to 0$, $\mu \to \infty$ and when $\eta \to \infty$, $\mu \to \lambda$. This implies a direct relationship between η and μ and in this sense, the two are inextricable.

since μ determines ω, it matters for equation (24) as well.

The second-order condition for equation (24) is $\mu - \lambda \geq 0$. When $\mu = \lambda$, the loss function is affine with respect to ω, meaning there is no uncertainty aversion. Were μ permitted to be less than λ, this would be tantamount to assuming the policymaker is an uncertainty seeker, in which case nature would always choose $\omega = -\infty$ guaranteeing the maximum possible loss to the authority. Thus the minimum value for μ is λ—that is, the H^{∞} solution—plus a small increment. When μ gets large, the solution approaches the LQG solution as we shall presently show.

Solving the first-order conditions for π and ω, we have:

$$\pi^r = \frac{-\lambda\gamma_0\gamma_1}{1 + \lambda\gamma_1^2 - \lambda/\mu} \tag{26}$$

$$\omega = \frac{\lambda\gamma_0}{\mu(1 + \lambda\gamma_1^2) - \lambda} \tag{27}$$

Notice that as the penalty on nature's control, μ, rises towards infinity, ω approaches zero, and $\pi \rightarrow -\lambda\gamma_0\gamma_1/(1 + \lambda\gamma_1^2)$ which is the certainty equivalent solution. Conversely, as $\mu \rightarrow \underline{\mu} = \lambda$, ω approaches $\gamma_0/\lambda\gamma_1^2$ from below and $\pi \rightarrow -\gamma_0/\gamma_1$.

It is easy to show that when the policymaker has doubts about the veracity of the reference model, inflation will be higher than in the certainty equivalent case:

$$\hat{\pi}^r - \hat{\pi}^{ce} = \frac{(\lambda\gamma_0\gamma_1)/\mu}{(1 + \lambda\gamma_1^2)(1 + \lambda\gamma_1^2 - \lambda/\mu)} > 0 \qquad \mu < \infty \tag{28}$$

The preceding has shown how unstructured model uncertainty and the policymaker's aversion to it can be modeled as a two-player game, and equivalently as a multiplier problem with the weight on the multiplier representing the authority's aversion to uncertainty. With this in mind, in Figures 4 and 5, we show the dynamic solutions for our model economy with two different settings for μ: $\mu = \lambda + 0.01$—which is arbitrarily close to the H^{∞} solution, the most highly uncertainty averse solution that is feasible in this framework—and $\mu = \lambda + 10$ which is mildly uncertainty aversive.

Figure 4
Inflation Performance and Perceived Supply Curve Parameters
(Robust response to unstructured model uncertainty; $\mu = 0.01$)

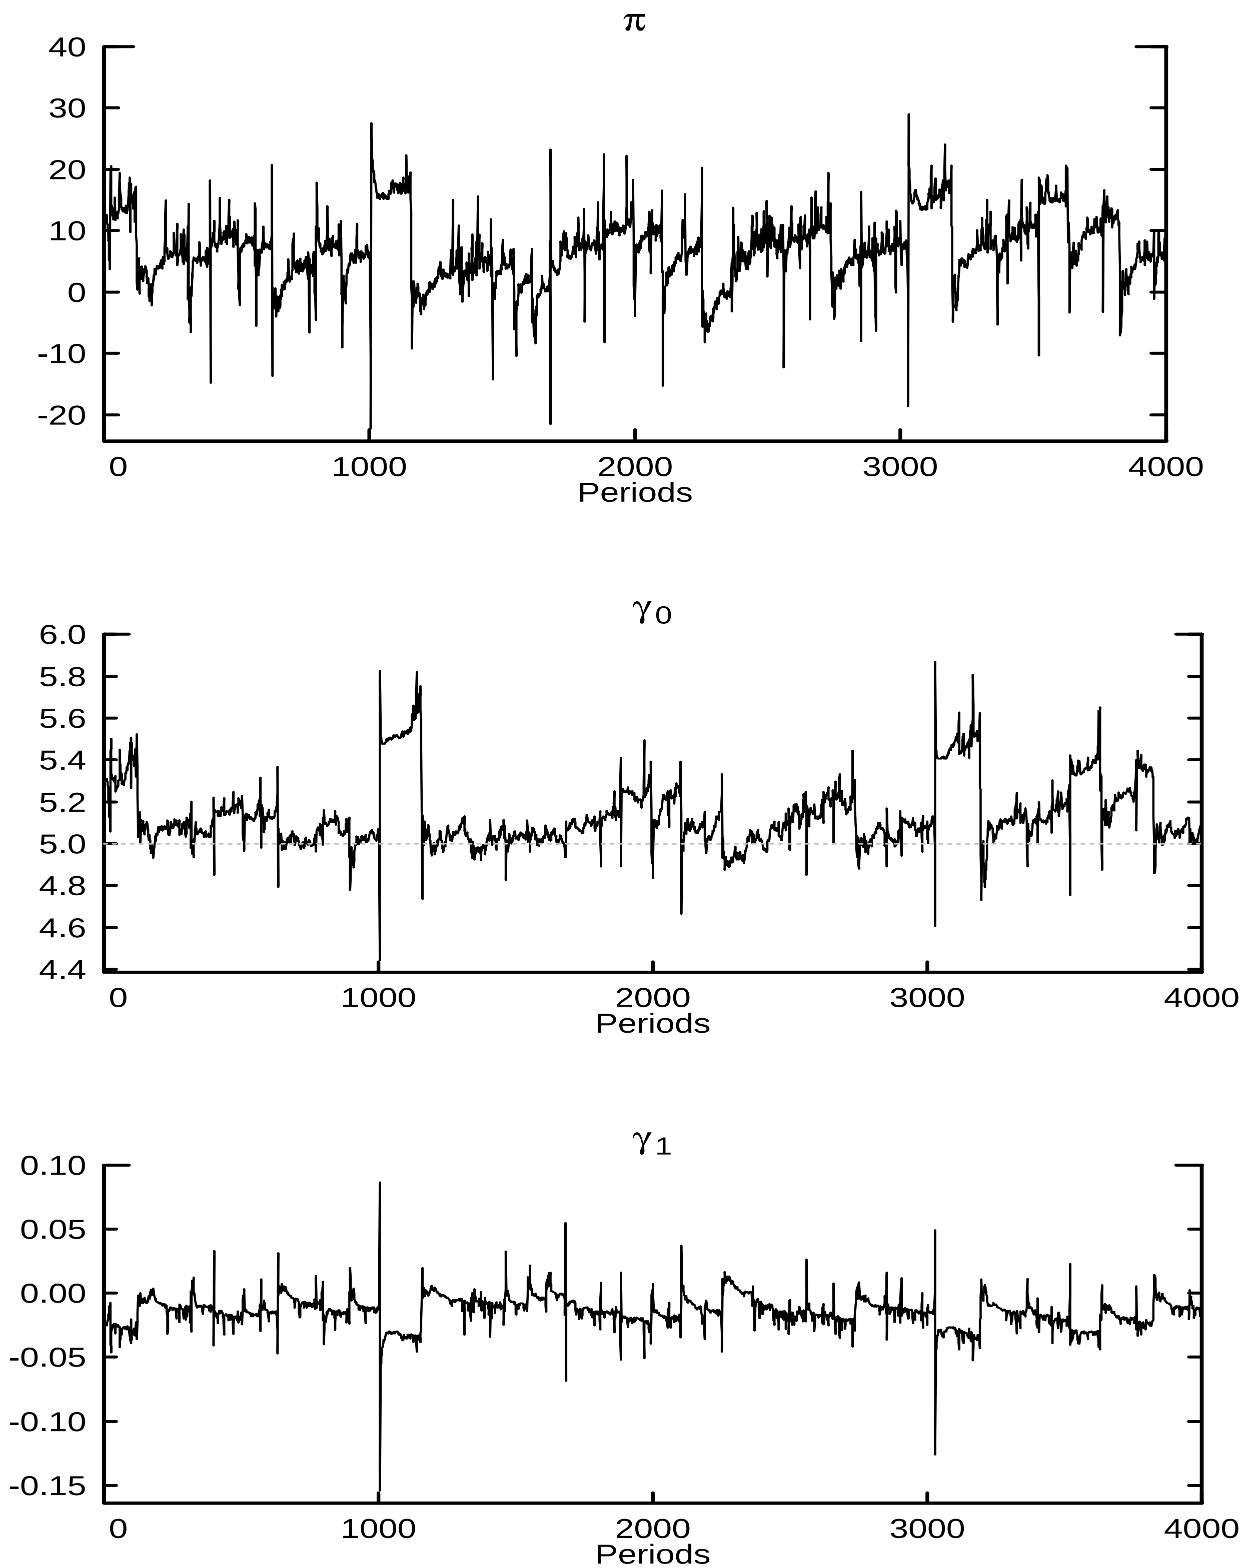

Figure 5
Inflation Performance and Perceived Supply Curve Parameters
(Robust response to unstructured model uncertainty; $\mu = 10$)

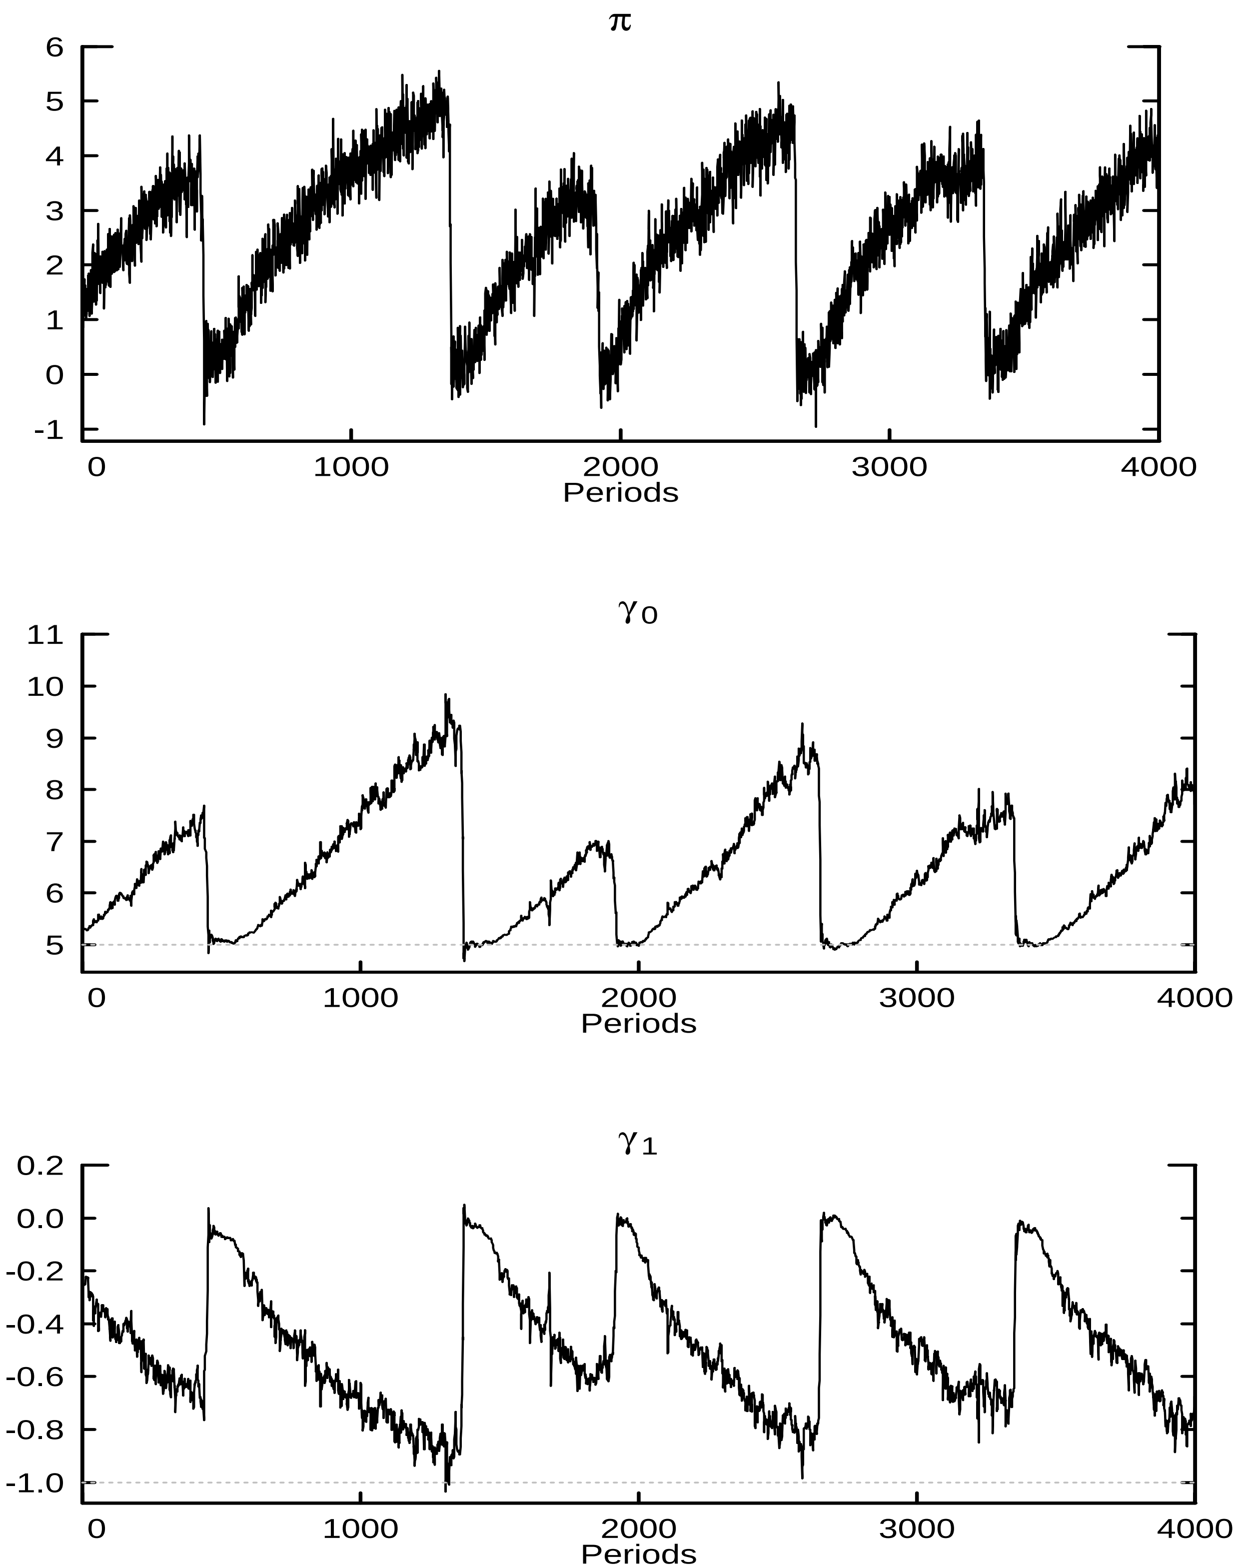

The solutions clearly show—particularly in comparison to Figure 1—that robust control as a response to the model uncertainty that is generated by the monetary authority's not understanding his or her own role in the data generating process only worsens economic performance with the degree of deterioration being a function of the degree of uncertainty aversion. Evidently, the extreme type of model uncertainty entertained here by the robust policymaker is not the right medicine for the comparatively mild specification error that is being committed here. From this observation we can also conclude that a policy maker facing Sargent's (1999) induction problem will not find the solution to that problem through the use of Hansen and Sargent's (1995) techniques.

3.5 welfare: a summing up

Figures 1 to 5 do a reasonable job of summarizing the performance of the economy. Most of what one needs to know to assess policy can be gleaned from the extremities of inflation that are reached in the simulations and the incidence of disinflations. Nevertheless, it is interesting to examine the extent to which welfare is affected, *ex post*, by the use of tools for handling model uncertainties.

To this end, we computed the losses from the stochastic simulations shown in the figures using the *ex post* calculation of equation (7). Eleven thousand observations were computed with the first 1000 discarded; the same set of stochastic shocks was used for each exercise. The results are shown in Table 1 below where the answer for the LQG rule has been normalized to unity without loss of generality. The table shows that all of the approaches that take uncertainty seriously nonetheless bring about no improvement in policy performance. (In point of fact, the Bayesian policymaker actually produces an improvement, but it is minuscule.) The robust control approaches to the induction problem are uniformly deleterious, with results that might be described as very poor when the authority is particularly uncertainty averse. This is true for both the unstructured model uncertainty approach, and the structured model uncertainty approach.

Table 1
Comparative loss from alternative approaches to model uncertainty

	Parameter settings	H^2 loss
Certainty equivalent	$\mu = \infty$	1.00
Bayesian	$\mu = \infty$	1.00
Unstructured robust	$\mu = 10$	1.03
Unstructured robust	$\mu = 0.01$	2.88
Structured robust	n/a	1.69

Notes: These are the results from stochastic simulations of the model for 11000 dates and discarding the first 1000 observations, evaluating equation (7).

However, for modest levels of (unstructured) uncertainty aversion, such as $\mu = 10$, the loss in performance is quite small suggesting that the small doses of robust policymaking to hedge against misspecifications *other than the one studied here* would not come at great cost. The fact remains, however, that for this popular rendition of model misspecification, the medicine of robust policy is worse than the disease. Taken at face value, these results suggest that the policymaker would be better off ignoring the advice of those who advocate the use of sophisticated tools to address uncertainty and practicing certainty equivalence instead.[14]

14. Qualitatively the same results obtain for much higher and much lower values of the preference parameter, λ, except that with very high λ, the loss with structured Knightian uncertainty can be slightly better than the certainty equivalent solution. As a specific example, when $\lambda = 10$, the loss under structured robust control is eight percent less than under LQG control.

4. Concluding remarks

This paper has re-examined Thomas Sargent's explanation of the great inflation of the 1970s. In the original work, the monetary authority was assumed to re-estimate his or her reference model in such a way as to suggest doubts about the constancy of the estimated parameters. And yet the policymaker was not allowed to carry forward those doubts to the question of policy design. We have relaxed this restriction by allowing the monetary authority to seek protection against model misspecification in three possible ways.

First, the policymaker was permitted to take the estimated standard errors of the parameters into account when designing policy, as a Bayesian would do. In this regard, we investigated the advice of Blinder (1998), based on Brainard (1967), among others. We found that the conventional advice to responding to model uncertainty—the Bayesian approach—produced results that were only infinitesimally different from the certainty equivalent case. In other words, if policy makers, operating in a world of uncertainty, were to follow Blinder's advice to "do less" than the certainty equivalent policy response to shocks, and did so using the usual Bayesian statistical criteria, they would find no solution to their problem.

Second, we allowed the policymaker to be a structured but robust controller, protecting against modeling errors of uncertain magnitude but known location. And third we allowed the policymaker to protect against unstructured modeling errors of unquantifiable location. In this case, we were following the advice of Sargent and Hansen (1995, 2001) on handling model uncertainty. In all instances, we found that protecting against model misspecification using this variety of sophisticated techniques made the policymaker worse off, rather than better off.

For the policymaker of today these results are not comforting. They suggest that there is no panacea for protecting against specification error and no substitute for the hard and often judgmental work of assessing whether models are "close enough" to be taken as given for rendering policy advice. At the same time, the results strengthen the case for the Sargent explanation of the inflation of the 1970s. Had the recurring bouts of Nash inflation followed by bouts of disinflation disappeared with the economy under the control of some or all of these policies, the results would have suggested that Sargent's findings were a manifestation of the assumed naiveté of the policy maker. Finally, there is some echoing in these results of other findings on where the gains accrue to a central bank's efforts at improving the conduct of monetary policy conduct. Orphanides *et al.* (2000) and Tetlow (2002) have both shown that the gains from designing policy to overcome an errors-in-variables problem are relatively small. Those findings suggest that investments in reducing the errors-in-variables problem by improving the quality of measurement may yield a larger return. The results described here support this line of argument.

This paper is the first contribution in this area of which we are aware. It would seem obvious that the results found here should be tested in more general frameworks with more sophisticated models. Looking beyond robust and Bayesian control ideas, the possibility also arises that robust filtering, or robust estimation, may produce reference models that are less susceptible to the induction problem (Hansen, Sargent and Tallarini (2000), Hansen, Sargent and Wang (2002)). This appears to us to be a promising direction in which to proceed.

5. References:

Barro, Robert J. and David B. Gordon (1983) “A positive theory of inflation in a natural rate model” 91,*Journal of Political Economy*,4 (August), pp. 589-610.

Blinder, Alan (1987) *Hard Heads, Soft Hearts: tough-minded economics for a just society* (Reading, MA: Addison-Wesley).

Blinder, Alan (1998) *Central Banking in Theory and Practice* (Cambridge, MA: MIT Press).

Brainard, William (1967) “Uncertainty and the Effectiveness of Monetary Policy” 57,*American Economic Review*,2 (May), pp. 411-425.

Bullard, James and In-Koo Cho (2001) “Escapist Policy Rules” unpublished manuscript, Federal Reserve Bank of St. Louis (December).

Caravani, P. (1995) “On H-infinity criteria for macroeconomic policy evaluation” 19,*Journal of Economic Dynamics and Control*,5/7 (July/September), pp. 961-984.

Cho, In-Koo; Noah Williams, and Thomas J. Sargent (2000) “Escaping Nash Inflation” unpublished manuscript, University of Chicago.

Chow, Gregory C. (1973) “Effect of Uncertainty on Optimal Control Policies” 14,*International Economic Review*,3 (October), pp. 632-645.

Cogley, Timothy and Thomas J. Sargent (2000) “Evolving Post-World War II U.S. Inflation Dynamics” unpublished manuscript (December).

Craine, Roger (1977) “Optimal Monetary Policy with Uncertainty” 1,*Journal of Economic Dynamics and Control*,1 (February), pp. 59-83.

De Long, J. Bradford (1997) “America’s peace-time inflation: the 1970s” in Christina D. Romer and David H. Romer (eds.) *Reducing Inflation: motivation and strategy* NBER Studies in Business Cycles, volume 30 (Chicago: University of Chicago Press), pp. 247-276.

Drew, Aaron and Benjamin Hunt (1999) “Efficient Simple Monetary Policy Rules and the Implications of Potential Output Uncertainty” Reserve Bank of New Zealand working paper no. G99/5 (August).

Giannoni, M. P. (2000) “Does model uncertainty justify caution?: robust optimal policy in a forward-looking model” forthcoming in *Macroeconomic Dynamics*.

Hall, Simon; Christopher Salmon; Tony Yates and Nicoletta Batini (1999) “Uncertainty and Simple Monetary Policy Rules: an illustration for the United Kingdom” Bank of England working paper no. 96 (June).

Hansen, Lars Peter and Thomas J. Sargent (1995) “Discounted linear exponential quadratic gaussian control” 40,*IEEE Transactions on Automatic Control*, pp. 968-971.

Hansen, Lars Peter and Thomas J. Sargent (2001) *Elements of Robust Control and Filtering for Macroeconomics* unpublished manuscript

Hansen, Lars Peter; Thomas J. Sargent, and T. Tallarini (1999) “Robust Permanent Income and Pricing” 66,*Review of Economic Studies*, pp. 873-907.

Hansen, Lars Peter; Thomas J. Sargent, and N. E. Wang (2002) “Robust Permanent Income and Pricing with Filtering” 6,*Macroeconomic Dynamics*,1 (February).

Ireland, Peter (1999) “Does the time-consistency problem explain the behavior of inflation in the United States?” 44,*Journal of Monetary Economics*,2 (October), pp. 279-292.

Kydland, Finn E. and Edward Prescott (1977) “Rules Rather than Discretion: the inconsistency of optimal plans” 85,*Journal of Political Economy* pp. 473-493.

Lucas, Robert E. (1972) “Econometric Testing of the Natural Rate Hypothesis” in Otto Eckstein (ed.) *Econometrics of Price Determination* (Washington: Board of Governors of the Federal Reserve System), pp. 50-59.

Lucas, Robert E. (1976) “Econometric Policy Evaluation: a critique” 1,*Carnegie-Rochester Conference Series on Public Policy*. pp. 19-46.

Martin, Ben (1999) “Caution and Gradualism in Monetary Policy Under Uncertainty” Bank of England working paper no. 125 (December).

Onatski, Alexei (1999) “Minimax Analysis of Model Uncertainty: comparison to Bayesian approach, worst possible economies, and optimal robust monetary policies” unpublished manuscript, Department of Economics, Harvard University.

Onatski, Alexei and James H. Stock (2002) “Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy” 6,*Macroeconomic Dynamics*,1 (February)

Orphanides, Athanasios; David Reifschneider, Robert Tetlow, John C. Williams and Frederico Finan (2000) “Errors in the Measurement of the Output Gap and the Design of Monetary Policy“ 52,*Journal of Economics and Business*,2 pp. 117-141.

Rudebusch, Glenn (2001) “Is the Fed Too Timid?: monetary policy in an uncertain world” 83,*Review of Economics and Statistics*,2 (May) pp. 203-217.

Sack, Brian (2000) “Does the Fed Act Gradually?: a VAR analysis” 46,*Journal of Monetary Economics*,1 pp. 229-256.

Sargent, Thomas J. (1971) “A Note on the ‘Accelerationist Controversy’” 3,*Journal of Money, Credit and Banking* pp. 721-725.

Sargent, Thomas J. (1999) *The Conquest of American Inflation* (Princeton: Princeton U. Press).

Schellekens, Phillip (1999) “Caution and Conservatism in the Making of Monetary Policy” unpublished manuscript, European Central Bank.

Shuetrim, Geoffrey and Christopher Thompson (2000) “The Implications of Uncertainty for Monetary Policy” in Benjamin Hunt and Adrian Orr (eds.) *Monetary Policy Under Uncertainty* (Wellington, NZ: Reserve Bank of New Zealand), pp. 259-292.

Smets, Frank (2000) “Output Gap Uncertainty: does it matter for the Taylor rule?” in Benjamin Hunt and Adrian Orr (eds.) *Monetary Policy Under Uncertainty* (Wellington, NZ: Reserve Bank of New Zealand), pp. 10-29.

Söderström, Ulf (1999) “Monetary Policy with Uncertain Parameters” working paper, Sveriges Riksbank, Stockholm.

Srour, Gabriel (1999) “Inflation Targeting Under Uncertainty” *Bank of Canada Technical Report no. 85* (Ottawa: Bank of Canada).

Svensson, Lars E.O. (2000) “Robust Control Made Simple” unpublished manuscript, Department of Economics, Princeton University.

Taylor, John B. (1997) “Comment on ‘America’s Only Peacetime Inflation: the 1970s’” in Christina D. Romer and David H. Romer (eds.) *Reducing Inflation: motivation and strategy* NBER Studies in Business Cycles, volume 30 (Chicago: University of Chicago Press).

Taylor, John B. (1998) “Monetary Policy Guidelines for Unemployment and Inflation Stability” in Robert M. Solow and John B. Taylor (eds.) *Inflation, Unemployment and Monetary Policy* (Cambridge, MA: MIT Press)

Tetlow, Robert J. (2002) “Uncertain Potential Output and Monetary Policy in a Forward-looking Model” unpublished manuscript in revision, Division of Research and Statistics, Board of Governors of the Federal Reserve System.

Tetlow, Robert J. and Peter von zur Muehlen (2001a) “Simplicity Versus Optimality: the choice monetary policy rules when agents must learn” 25,*Journal of Economic Dynamics and Control*,1/2 pp. 245-279.

Tetlow, Robert J. and Peter von zur Muehlen (2001b) “Robust Monetary Policy with Misspecified Models: Does model uncertainty always call for attenuated policy? 25,*Journal of Economic Dynamics and Control*,6/7 pp. 911-949.

Tinbergen, Jan (1952) *On the Theory of Economic Policy*, second edition, volume I of Contributions to Economic Analysis (Amsterdam: North Holland).

Theil, Henri (1961) *Economic Forecasts and Policy*, second edition, volume XV of Contributions to Economic Analysis (Amsterdam: North Holland).

von zur Muehlen, Peter (1982) “Activist versus Nonactivist Monetary Policy: optimal rules under extreme uncertainty” manuscript now available as Finance and Economics Discussion Series paper no. 2001-02 (January).

Whittle, Peter (1990) *Risk-sensitive Optimal Control* (New York: Wiley).

Williams, Noah (2000) “Escape Dynamics in Learning Models” unpublished manuscript, University of Chicago (November).